

03-31-00
A
IN THE UNITED STATES
PATENT AND TRADEMARK OFFICE

PATENT APPLICATION

"Express Mail" Label No. : EL525102993US

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 C.F.R. §1.10 on the date indicated below and is addressed to: Commissioner of Patents and Trademarks, Washington, D.C. 20231

Date of Deposit: March 30, 2000

Signature: Linda M. Shuckleton

INVENTOR(S) Alan David Berenbaum,
Nevin Heintze,
Tor E. Jeremiassen and
Stefanos Kaxiras

Berenbaum 9-4-5-5

TITLE METHOD AND APPARATUS FOR SPLITTING
PACKETS IN A MULTITHREADED VLIW PROCESSOR

ASSISTANT COMMISSIONER FOR PATENTS
WASHINGTON, D.C. 20231

SIR:

Enclosed are the following papers relating to the above-named application for patent:

Specification
6 Sheets of informal drawing(s)
Declaration and Power of Attorney
1 Assignment with Cover Sheet

CLAIMS AS FILED				
	NO. FILED	NO. EXTRA	RATE	CALCULATIONS
Total Claims	16 - 20 =	0	X \$18 =	\$ 0
Independent Claims	4 - 3 =	1	X \$78 =	\$ 78
Multiple Dependent Claim(s), if applicable			\$260 =	\$ 0
Basic Fee				\$690
			TOTAL FEE:	\$768

Please file the application and charge **Lucent Technologies Deposit Account No. 12-2325** the amount of \$768.00, to cover the filing fee. *Duplicate copies of this letter are enclosed.* In the event of non-payment or improper payment of a required fee, the Commissioner is authorized to charge or to credit **Deposit Account No. 12-2325** as required to correct the error.

Please address all correspondence to: Kevin M. Mason, Ryan & Mason, L.L.P., 90 Forest Avenue, Locust Valley, New York 11560. Telephone calls should be made to the under-signed attorney at (203) 255-6560.

Respectfully,

Kevin M. Mason

Date: March 30, 2000

Kevin M. Mason
Reg. No. 36,597
Attorney for Applicant(s)

Ryan & Mason, L.L.P.
90 Forest Avenue
Locust Valley, New York 11560

METHOD AND APPARATUS FOR SPLITTING PACKETS
IN A MULTITHREADED VLIW PROCESSOR

Cross-Reference to Related Applications

5 The present invention is related to United States Patent Application entitled
“Method and Apparatus for Allocating Functional Units in a Multithreaded Very Large
Instruction Word (VLIW) Processor,” Attorney Docket Number (Berenbaum 7-2-3-3); United
States Patent Application entitled “Method and Apparatus for Releasing Functional Units in a
Multithreaded Very Large Instruction Word (VLIW) Processor,” Attorney Docket Number
10 (Berenbaum 8-3-4-4); and United States Patent Application entitled “Method and Apparatus for
Identifying Splittable Packets in a Multithreaded Very Large Instruction Word (VLIW)
Processor,” Attorney Docket Number (Berenbaum 10-5-6-6), each filed contemporaneously
herewith, assigned to the assignee of the present invention and incorporated by reference herein.

15 **Field of the Invention**

The present invention relates generally to multithreaded processors, and, more
particularly, to a method and apparatus for splitting packets in such multithreaded processors.

20 **Background of the Invention**

Computer architecture designs attempt to complete workloads more quickly. A
number of architecture designs have been proposed or suggested for exploiting program
parallelism. Generally, an architecture that can issue more than one operation at a time is capable
of executing a program faster than an architecture that can only issue one operation at a time.
Most recent advances in computer architecture have been directed towards methods of issuing
25 more than one operation at a time and thereby speed up the operation of programs. FIG. 1
illustrates a conventional microprocessor architecture 100. Specifically, the microprocessor 100
includes a program counter (PC) 110, a register set 120 and a number of functional units (FUs)
130-N. The redundant functional units (FUs) 130-1 through 130-N provide the illustrative
microprocessor architecture 100 with sufficient hardware resources to perform a corresponding
30 number of operations in parallel.

An architecture that exploits parallelism in a program issues operands to more than one functional unit at a time to speed up the program execution. A number of architectures have been proposed or suggested with a parallel architecture, including superscalar processors, very long instruction word (VLIW) processors and multithreaded processors, each discussed below in conjunction with FIGS. 2, 4 and 5, respectively. Generally, a superscalar processor utilizes hardware at run-time to dynamically determine if a number of operations from a single instruction stream are independent, and if so, the processor executes the instructions using parallel arithmetic and logic units (ALUs). Two instructions are said to be independent if none of the source operands are dependent on the destination operands of any instruction that precedes them. A very long instruction word (VLIW) processor evaluates the instructions during compilation and groups the operations appropriately, for parallel execution, based on dependency information. A multithreaded processor, on the other hand, executes more than one instruction stream in parallel, rather than attempting to exploit parallelism within a single instruction stream.

A superscalar processor architecture 200, shown in FIG. 2, has a number of functional units that operate independently, in the event each is provided with valid data. For example, as shown in FIG. 2, the superscalar processor 200 has three functional units embodied as arithmetic and logic units (ALUs) 230-N, each of which can compute a result at the same time. The superscalar processor 200 includes a front-end section 208 having an instruction fetch block 210, an instruction decode block 215, and an instruction sequencing unit 220 (issue block). The instruction fetch block 210 obtains instructions from an input queue 205 of a single threaded instruction stream. The instruction sequencing unit 220 identifies independent instructions that can be executed simultaneously in the available arithmetic and logic units (ALUs) 230-N, in a known manner. The refine block 250 allows the instructions to complete, and also provides buffering and reordering for writing results back to the register set 240.

In the program fragment 310 shown in FIG. 3, instructions in locations L1, L2 and L3 are independent, in that none of the source operands in instructions L2 and L3 are dependent on the destination operands of any instruction that precedes them. When the program counter (PC) is set to location L1, the instruction sequencing unit 220 will look ahead in the instruction stream and detect that the instructions at L2 and L3 are independent, and thus all three can be issued simultaneously to the three available functional units 230-N. For a more detailed

discussion of superscalar processors, see, for example, James. E. Smith and Gurindar. S. Sohi, "The Microarchitecture of Superscalar Processors," Proc. of the IEEE (Dec. 1995), incorporated by reference herein.

As previously indicated, a very long instruction word (VLIW) processor 400, shown in FIG.4, relies on software to detect data parallelism at compile time from a single instruction stream, rather than using hardware to dynamically detect parallelism at run time. A VLIW compiler, when presented with the source code that was used to generate the code fragment 310 in FIG. 3, would detect the instruction independence and construct a single, very long instruction comprised of all three operations. At run time, the issue logic of the processor 400 would issue this wide instruction in one cycle, directing data to all available functional units 430-N. As shown in FIG. 4, the very long instruction word (VLIW) processor 400 includes an integrated fetch/decode block 420 that obtains the previously grouped instructions 410 from memory. For a more detailed discussion of very long instruction word (VLIW) processors, see, for example, Burton J. Smith, "Architecture and Applications of the HEP Multiprocessor Computer System," SPIE Real Time Signal Processing IV, 241-248 (1981), incorporated by reference herein.

One variety of VLIW processors, for example, represented by the Multiflow architecture, discussed in Robert P. Colwell et al., "A VLIW Architecture for a Trace Scheduling Compiler," IEEE Transactions on Computers (August 1988), uses a fixed-width instruction, in which predefined fields direct data to all functional units 430-N at once. When all operations specified in the wide instruction are completed, the processor issues a new, multi-operation instruction. Some more recent VLIW processors, such as the C6x processor commercially available from Texas Instruments, of Dallas, TX and the EPIC IA-64 processor commercially available from Intel Corp, of Santa Clara, CA, instead use a variable-length instruction packet, which contains one or more operations bundled together.

A multithreaded processor 500, shown in FIG. 5, gains performance improvements by executing more than one instruction stream in parallel, rather than attempting to exploit parallelism within a single instruction stream. The multithreaded processor 500 shown in FIG. 5 includes a program counter 510-N, a register set 520-N and a functional unit 530-N, each dedicated to a corresponding instruction stream N. Alternate implementations of the

multithreaded processor 500 have utilized a single functional unit 530, with several register sets 520-N and program counters 510-N. Such alternate multithreaded processors 500 are designed in such a way that the processor 500 can switch instruction issue from one program counter/register set 510-N/520-N to another program counter/register set 510-N/520-N in one or two cycles. A long latency instruction, such as a LOAD instruction, can thus be overlapped with shorter operations from other instruction streams. The TERA MTA architecture, commercially available from Tera Computer Company, of Seattle, WA, is an example of this type.

An extension of the multithreaded architecture 500, referred to as Simultaneous Multithreading, combines the superscalar architecture, discussed above in conjunction with FIG. 2, with the multithreaded designs, discussed above in conjunction with FIG. 5. For a detailed discussion of Simultaneous Multithreading techniques, see, for example, Dean Tullsen et al., "Simultaneous Multithreading: Maximizing On-Chip Parallelism," Proc. of the 22nd Annual Int'l Symposium on Computer Architecture, 392-403 (Santa Margherita Ligure, Italy, June 1995), incorporated by reference herein. Generally, in a Simultaneous Multithreading architecture, there is a pool of functional units, any number of which may be dynamically assigned to an instruction which can issue from any one of a number of program counter/register set structures. By sharing the functional units among a number of program threads, the Simultaneous Multithreading architecture can make more efficient use of hardware than that shown in FIG. 5.

While the combined approach of the Simultaneous Multithreading architecture provides improved efficiency over the individual approaches of the superscalar architecture or the multithreaded architecture, Simultaneous Multithreading architectures still require elaborate issue logic to dynamically examine instruction streams in order to detect potential parallelism. In addition, when an operation takes multiple cycles, the instruction issue logic may stall, because there is no other source of operations available. Conventional multithreaded processors issue instructions from a set of instructions simultaneously, with the functional units designed to accommodate the widest potential issue. A need therefore exists for a multithreaded processor architecture that does not require a dynamic determination of whether or not two instruction streams are independent. A further need exists for a multithreaded architecture that provides

simultaneous multithreading. Yet another need exists for a method and apparatus that improve the utilization of processor resources for each cycle.

Summary of the Invention

5 Generally, a method and apparatus are disclosed for processing instructions in a multithreaded very large instruction word (VLIW) processor. The present invention combines the techniques of conventional very long instruction word (VLIW) architectures and conventional multithreaded architectures. The combined architecture of the present invention reduces execution time within an individual program, as well as across a workload. The present
10 invention utilizes instruction packet splitting to recover some efficiency lost with conventional multithreaded architectures. Instruction packet splitting allows an instruction bundle to be partially issued in one cycle, with the remainder of the bundle issued during a subsequent cycle. Thus, the present invention provides greater utilization of hardware resources (such as the functional units) and a lower elapsed time across a workload comprising multiple threads.

15 Instruction packet splitting increases throughput across all instruction threads, and reduces the number of cycles that the functional units rest idle. The allocation hardware of the present invention assigns as many instructions from each packet as will fit on the available functional units, rather than allocating all instructions in an instruction packet at one time. Those instructions that cannot be allocated to a functional unit are retained in a ready-to-run register.
20 On subsequent cycles, instruction packets in which all instructions have been issued to functional units are updated from their thread's instruction stream, while instruction packets with instructions that have been held are retained. The functional unit allocation logic can then assign instructions from the newly-loaded instruction packets as well as instructions that were not issued from the retained instruction packets.

25 The present invention utilizes a compiler to detect parallelism in a multithreaded processor architecture. Thus, a multithreaded VLIW architecture is disclosed that exploits program parallelism by issuing multiple instructions, in a similar manner to single threaded VLIW processors, from a single program sequencer, and also supporting multiple program sequencers, as in simultaneous multithreading but with reduced complexity in the issue logic,
30 since a dynamic determination is not required.

A more complete understanding of the present invention, as well as further features and advantages of the present invention, will be obtained by reference to the following detailed description and drawings.

Brief Description of the Drawings

FIG. 1 illustrates a conventional generalized microprocessor architecture;

FIG. 2 is a schematic block diagram of a conventional superscalar processor architecture;

FIG. 3 is a program fragment illustrating the independence of operations;

FIG. 4 is a schematic block diagram of a conventional very long instruction word (VLIW) processor architecture;

FIG. 5 is a schematic block diagram of a conventional multithreaded processor;

FIG. 6 illustrates a multithreaded VLIW processor in accordance with the present invention;

FIG. 7A illustrates a conventional pipeline for a multithreaded processor;

FIG. 7B illustrates a pipeline for a multithreaded processor in accordance with the present invention;

FIG. 8 is a schematic block diagram of an implementation of the allocate stage of FIG. 7B;

FIG. 9 illustrates the execution of three threads TA-TC for a conventional multithreaded implementation, where threads B and C have higher priority than thread A; and

FIGS. 10A and 10B illustrate the operation of instruction packet splitting in accordance with the present invention.

Detailed Description

FIG. 6 illustrates a Multithreaded VLIW processor 600 in accordance with the present invention. As shown in FIG. 6, there are three instruction threads, namely, thread A (TA), thread B (TB) and thread C (TC), each operating at instruction number n . In addition, the illustrative Multithreaded VLIW processor 600 includes nine functional units 620-1 through 620-9, which can be allocated independently to any thread TA-TC. Since the number of instructions

across the illustrative three threads TA-TC is nine and the illustrative number of available functional units 620 is also nine, then each of the instructions from all three threads TA-TC can issue their instruction packets in one cycle and move onto instruction $n+1$ on the subsequent cycle.

It is noted that there is generally a one-to-one correspondence between instructions and the operation specified thereby. Thus, such terms are used interchangeably herein. It is further noted that in the situation where an instruction specifies multiple operations, it is assumed that the multithreaded VLIW processor 600 includes one or more multiple-operation functional units 620 to execute the instruction specifying multiple operations. An example of an architecture where instructions specifying multiple operations may be processed is a complex instruction set computer (CISC).

The present invention allocates instructions to functional units to issue multiple VLIW instructions to multiple functional units in the same cycle. The allocation mechanism of the present invention occupies a pipeline stage just before arguments are dispatched to functional units. Thus, FIG. 7A illustrates a conventional pipeline 700 comprised of a fetch stage 710, where a packet is obtained from memory, a decode stage 720, where the required functional units and registers are identified for the fetched instructions, and an execute stage 730, where the specified operations are performed and the results are processed.

Thus, in a conventional VLIW architecture, a packet containing up to K instructions is fetched each cycle (in the fetch stage 710). Up to K instructions are decoded in the decode stage 720 and sent to (up to) K functional units (FUs). The registers corresponding to the instructions are read, the functional units operate on them and the results are written back to registers in the execute stage 730. It is assumed that up to three registers can be read and up to one register can be written per functional unit.

FIG. 7B illustrates a pipeline 750 in accordance with the present invention, where an allocate stage 780, discussed further below in conjunction with FIG. 8, is added for the implementation of multithreaded VLIW processors. Generally, the allocate stage 780 determines how to group the operations together to maximize efficiency. Thus, the pipeline 750 includes a fetch stage 760, where up to N packets are obtained from memory, a decode stage 770, where the functional units and registers are identified for the fetched instructions (up to $N*K$ instructions),

an allocate stage 780, where the appropriate instructions are selected and assigned to the FUs, and an execute stage 790, where the specified operations are performed and the results are processed.

In the multithreaded VLIW processor 600 of the present invention, up to N threads are supported in hardware. N thread contexts exist and contain all possible registers of a single thread and all status information required. A multithreaded VLIW processor 600 has M functional units, where M is greater than or equal to K. The modified pipeline stage 750, shown in FIG. 7B, works in the following manner. In each cycle, up to N packets (each containing up to K instructions) are fetched at the fetch stage 760. The decode stage 770 decodes up to N*K instructions and determines their requirements and the registers read and written. The allocate stage 780 selects M out of (up to) N*K instructions and forwards them to the M functional units. It is assumed that each functional unit can read up to 3 registers and write one register. In the execute stage 790, up to M functional units read up to 3*M registers and write up to M registers.

The allocate stage 780 selects for execution the appropriate M instructions from the (up to) N * K instructions that were fetched and decoded at stages 760 and 770. The criteria for selection are thread priority or resource availability or both. Under the thread priority criteria, different threads can have different priorities. The allocate stage 780 selects and forwards the packets (or instructions from packets) for execution belonging to the thread with the highest priority according to the priority policy implemented. A multitude of priority policies can be implemented. For example, a priority policy for a multithreaded VLIW processor supporting N contexts (N hardware threads) can have N priority levels. The highest priority thread in the processor is allocated before any other thread. Among threads with equal priority, the thread that waited the longest for allocation is preferred.

Under the resource availability criteria, a packet (having up to K instructions) can be allocated only if the resources (functional units) required by the packet are available for the next cycle. Functional units report their availability to the allocate stage 780.

FIG. 8 illustrates a schematic block diagram of an implementation of the allocate stage 780. As shown in FIG. 8, the hardware needed to implement the allocate stage 780 includes a priority encoder 810 and two crossbar switches 820, 830. Generally, the priority encoder 810 examines the state of the multiple operations in each thread, as well as the state of

the available functional units. The priority encoder 810 selects the packets that are going to execute and sets up the first crossbar switch 820 so that the appropriate register contents are transferred to the functional units at the beginning of the next cycle. The output of the priority encoder 810 configures the first crossbar switch 820 to route data from selected threads to the appropriate functional units. This can be accomplished, for example, by sending the register identifiers (that include a thread identifier) to the functional units and letting the functional units read the register contents via a separate data network and using the crossbar switch 810 to move the appropriate register contents to latches that are read by the functional units at the beginning of the next cycle

Out of the N packets that are fetched by the fetch stage 760 (FIG. 7B), the priority encoder 810 selects up to N packets for execution according to priority and resource availability. In other words, the priority encoder selects the highest priority threads that do not request unavailable resources for execution. It then sets up the first crossbar switch 810. The input crossbar switch 810 routes up to $3K*N$ inputs to up to $3*M$ outputs. The first crossbar switch 810 has the ability to transfer the register identifiers (or the contents of the appropriate registers) of each packet to the appropriate functional units.

Since there are up to N threads that can be selected in the same cycle and each thread can issue a packet of up to K instructions and each instruction can read up to 3 registers there are $3K*N$ register identifiers to select from. Since there are only M functional units and each functional unit can accept a single instruction, there are only $3M$ register identifiers to be selected. Therefore, the crossbar switch implements a $3K*N$ to $3M$ routing of register identifiers (or register contents).

The output crossbar switch 830 routes M inputs to $N*M$ or $N*K$ outputs. The second crossbar switch 830 is set up at the appropriate time to transfer the results of the functional units back to the appropriate registers. The second crossbar switch 830 can be implemented as a separate network by sending the register identifiers (that contain a thread identifier) to the functional units. When a functional unit computes a result, the functional unit routes the result to the given register identifier. There are M results that have to be routed to up to N threads. Each thread can accept up to K results. The second crossbar switch 830 routes M results to $N*K$ possible destinations. The second crossbar switch 830 can be implemented as M

buses that are connected to all N register files. In this case, the routing becomes M results to $N*M$ possible destinations (if the register files have the ability to accept M results).

In a conventional single-threaded VLIW architecture, all operations in an instruction packet are issued simultaneously. There are always enough functional units available to issue a packet. When an operation takes multiple cycles, the instruction issue logic may stall, because there is no other source of operations available. In a multithreaded VLIW processor in accordance with the present invention, on the other hand, these restrictions do not apply.

FIG. 9 illustrates the execution of three threads TA-TC for a conventional multithreaded implementation (without the benefit of the present invention), where threads B and C have higher priority than thread A. Since thread A runs at the lowest priority, its operations will be the last to be assigned. As shown in FIG. 9, five functional units 920 are assigned to implement the five operations in the current cycle of the higher priority threads TB and TC. Thread A has four operations, but there are only two functional units 920 available. Thus, thread A stalls for a conventional multithreaded implementation.

In order to maximize throughput across all threads, and minimize the number of cycles that the functional units rest idle, the present invention utilizes instruction packet splitting. Instead of allocating all operations in an instruction packet at one time, the allocation hardware 780, discussed above in conjunction with FIG. 8, assigns as many operations from each packet as will fit on the available functional units. Those operations that will not fit are retained in a ready-to-run register 850 (FIG. 8). On subsequent cycles, instruction packets in which all operations have been issued to functional units are updated from their thread's instruction stream, while instruction packets with operations that have been held are retained. The functional unit allocation logic 780 can then assign operations from the newly-loaded instruction packets as well as operations that were not issued from the retained instruction packets.

The operation of instruction packet splitting in accordance with the present invention is illustrated in FIGS. 10A and 10B. In FIG. 10A, there are three threads, each with an instruction packet from location n ready to run at the start of cycle x . Thread A runs at the lowest priority, so its operations will be the last to be assigned. Threads B and C require five of the seven available functional units 1020 to execute. Only two functional units 1020-2 and 1020-6

remain, so only the first two operations from thread A are assigned to execute. All seven functional units 1020 are now fully allocated.

At the completion of cycle x , the instruction packets for threads B and C are retired. The instruction issue logic associated with the threads replaces the instruction packets with those for address $n+1$, as shown in FIG. 10B. Since the instruction packet for thread A is not completed, the packet from address n is retained, with the first two operations marked completed. On the next cycle, $x+1$, illustrated in FIG. 10B, the final two operations from thread A are allocated to functional units, as well as all the operations from threads B and C.

Thus, the present invention provides greater utilization of hardware resources (i.e., the functional units 1020) and a lower elapsed time across a workload comprising the multiple threads.

It is to be understood that the embodiments and variations shown and described herein are merely illustrative of the principles of this invention and that various modifications may be implemented by those skilled in the art without departing from the scope and spirit of the invention.

We claim:

1. A multithreaded very large instruction word (VLIW) processor, comprising:
a plurality of functional units for executing instructions from a multithreaded instruction stream, said instructions being grouped into instruction packets by a compiler; and
an allocator that selects instructions from said instruction stream and forwards said instructions to said plurality of functional units, said allocator assigning instructions from at least one of said instruction packets to a plurality of said functional units.
2. The multithreaded very large instruction word (VLIW) processor of claim 1, wherein said allocator assigns as many instructions from a given instruction packet as permitted by an availability of said functional units.
3. The multithreaded very large instruction word (VLIW) processor of claim 1, further comprising a register for storing for execution in a later cycle an indication of those instructions from a given instruction packet that cannot be allocated to a functional unit in a given cycle.
4. The multithreaded very large instruction word (VLIW) processor of claim 3, wherein instruction packets in which all instructions have been issued to functional units are updated from the instruction stream of said thread.
5. The multithreaded very large instruction word (VLIW) processor of claim 3, wherein instruction packets with instructions indicated in said register are retained.
6. The multithreaded very large instruction word (VLIW) processor of claim 1, wherein said allocator can split an instruction packet provided the semantics of the instruction packet assembled by the compiler are not violated.
7. The multithreaded very large instruction word (VLIW) processor of claim 1, wherein said allocator can split an instruction packet provided a source register for one of the instructions

in a first part of said packet is not modified by one of the instructions in a second part of said packet.

8. A method for processing instructions in a multithreaded very large instruction word (VLIW) processor, comprising the steps of:

executing said instructions using a plurality of functional units, wherein said instructions are grouped into instruction packets by a compiler;

assigning instructions from said instruction stream to said plurality of functional units, wherein instructions from at least one of said instruction packets are assigned to a plurality of said functional units; and

forwarding said selected instructions to said corresponding functional units.

9. The method of claim 8, wherein said assigning step assigns as many instructions from a given instruction packet as permitted by an availability of said functional units.

10. The method of claim 8, further comprising the step of storing for execution in a later cycle an indication of those instructions from a given instruction packet that cannot be allocated to a functional unit in a given cycle.

11. The method of claim 10, wherein instruction packets in which all instructions have been issued to functional units are updated from the instruction stream of said thread.

12. The method of claim 10, wherein instruction packets with instructions indicated in said register are retained.

13. The method of claim 8, wherein said assigning step can split an instruction packet provided the semantics of the instruction packet assembled by the compiler are not violated.

14. The method of claim 8, wherein said assigning step can split an instruction packet provided a source register for one of the instructions in a first part of said packet is not modified by one of the instructions in a second part of said packet.

5 15. An article of manufacture for processing instructions from an instruction stream having a plurality of threads in a multithreaded very large instruction word (VLIW) processor, comprising:

a computer readable medium having computer readable program code means embodied thereon, said computer readable program code means comprising program code means
10 for causing a computer to:

execute said instructions using a plurality of functional units, wherein said instructions are grouped into instruction packets by a compiler;

assign instructions from said instruction stream to said plurality of functional units, wherein instructions from at least one of said instruction packets are assigned to a plurality
15 of said functional units; and

forward said selected instructions to said corresponding functional units.

16. A multithreaded very large instruction word (VLIW) processor, comprising:


a plurality of functional units for executing instructions from a multithreaded
20 instruction stream, said instructions being grouped into instruction packets by a compiler; and

an allocator that selects instructions from said instruction stream and forwards said instructions to said plurality of functional units, said allocator assigning instructions from at least one of said instruction packets to a plurality of said functional units provided the semantics of said instruction packet assembled by said compiler are not violated.

ABSTRACT


A method and apparatus are disclosed for allocating functional units in a multithreaded very large instruction word (VLIW) processor. The present invention combines the techniques of conventional very long instruction word (VLIW) architectures and conventional multithreaded architectures to reduce execution time within an individual program, as well as across a workload. The present invention utilizes instruction packet splitting to recover some efficiency lost with conventional multithreaded architectures. Instruction packet splitting allows an instruction bundle to be partially issued in one cycle, with the remainder of the bundle issued during a subsequent cycle. The allocation hardware assigns as many instructions from each packet as will fit on the available functional units, rather than allocating all instructions in an instruction packet at one time. Those instructions that cannot be allocated to a functional unit are retained in a ready-to-run register. On subsequent cycles, instruction packets in which all instructions have been issued to functional units are updated from their thread's instruction stream, while instruction packets with instructions that have been held are retained. The functional unit allocation logic can then assign instructions from the newly-loaded instruction packets as well as instructions that were not issued from the retained instruction packets.

1200-407.app



PRIOR ART

FIG. 1



PRIOR ART

FIG. 2

310 { L3: ADD R0, R1, R2
L2: SUB R3, R4, R2
L1: OR R6, R1, R5

FIG. 3

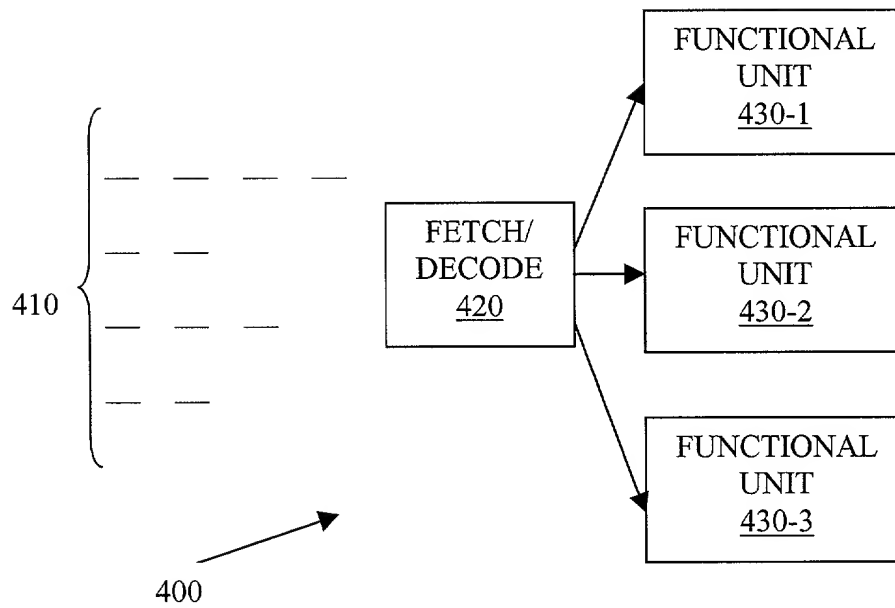


FIG. 4

PRIOR ART

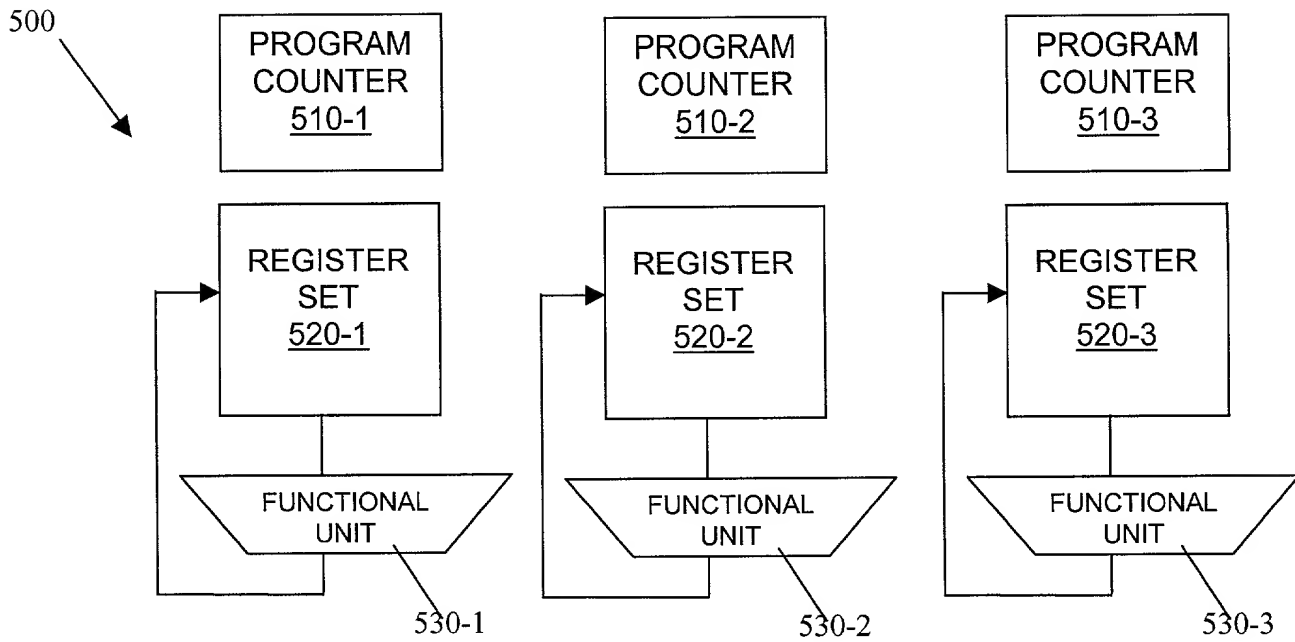


FIG. 5

PRIOR ART

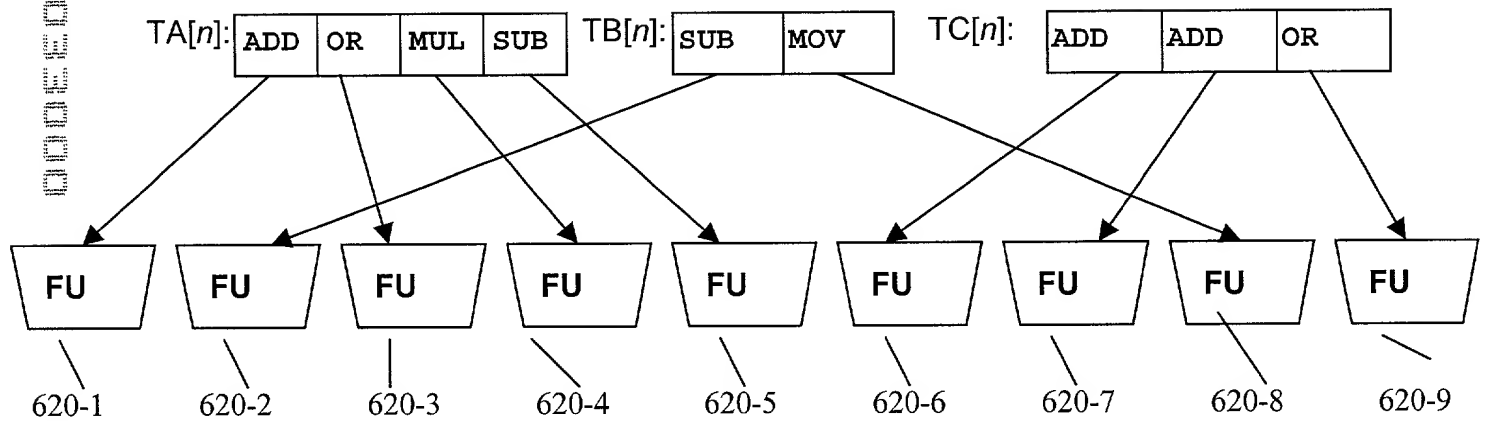


FIG. 6

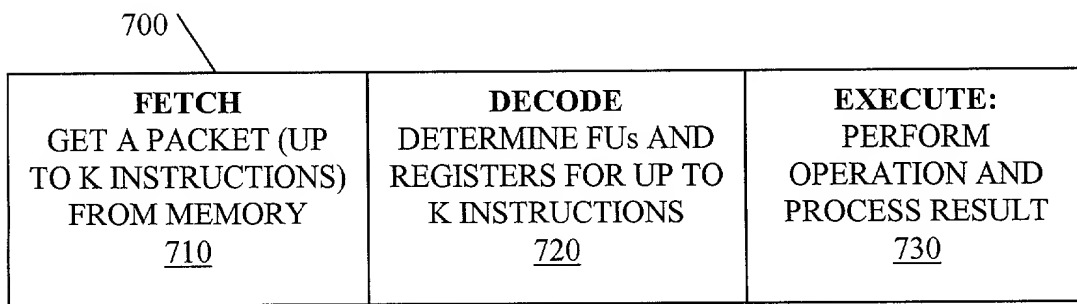


FIG. 7A **PRIOR ART**

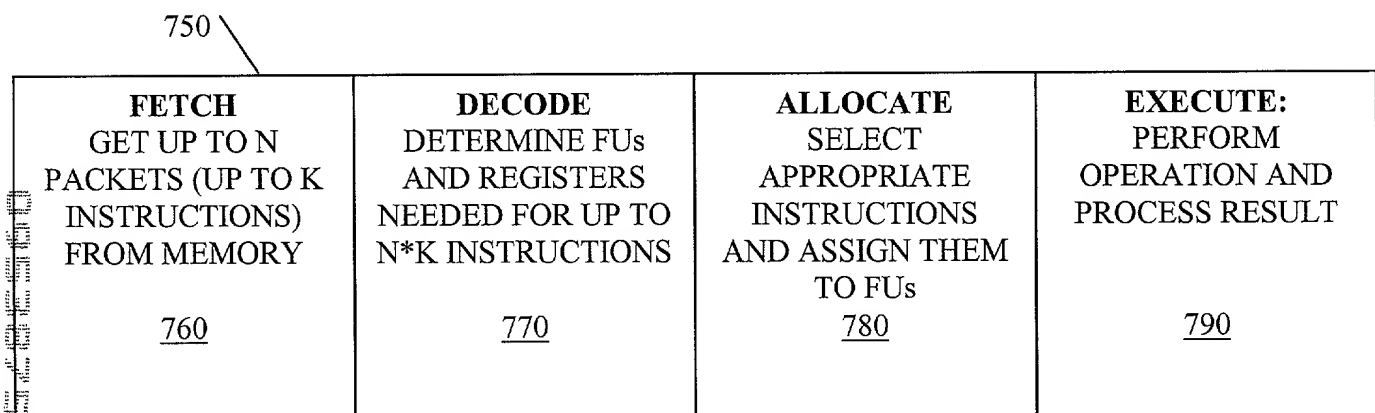


FIG. 7B

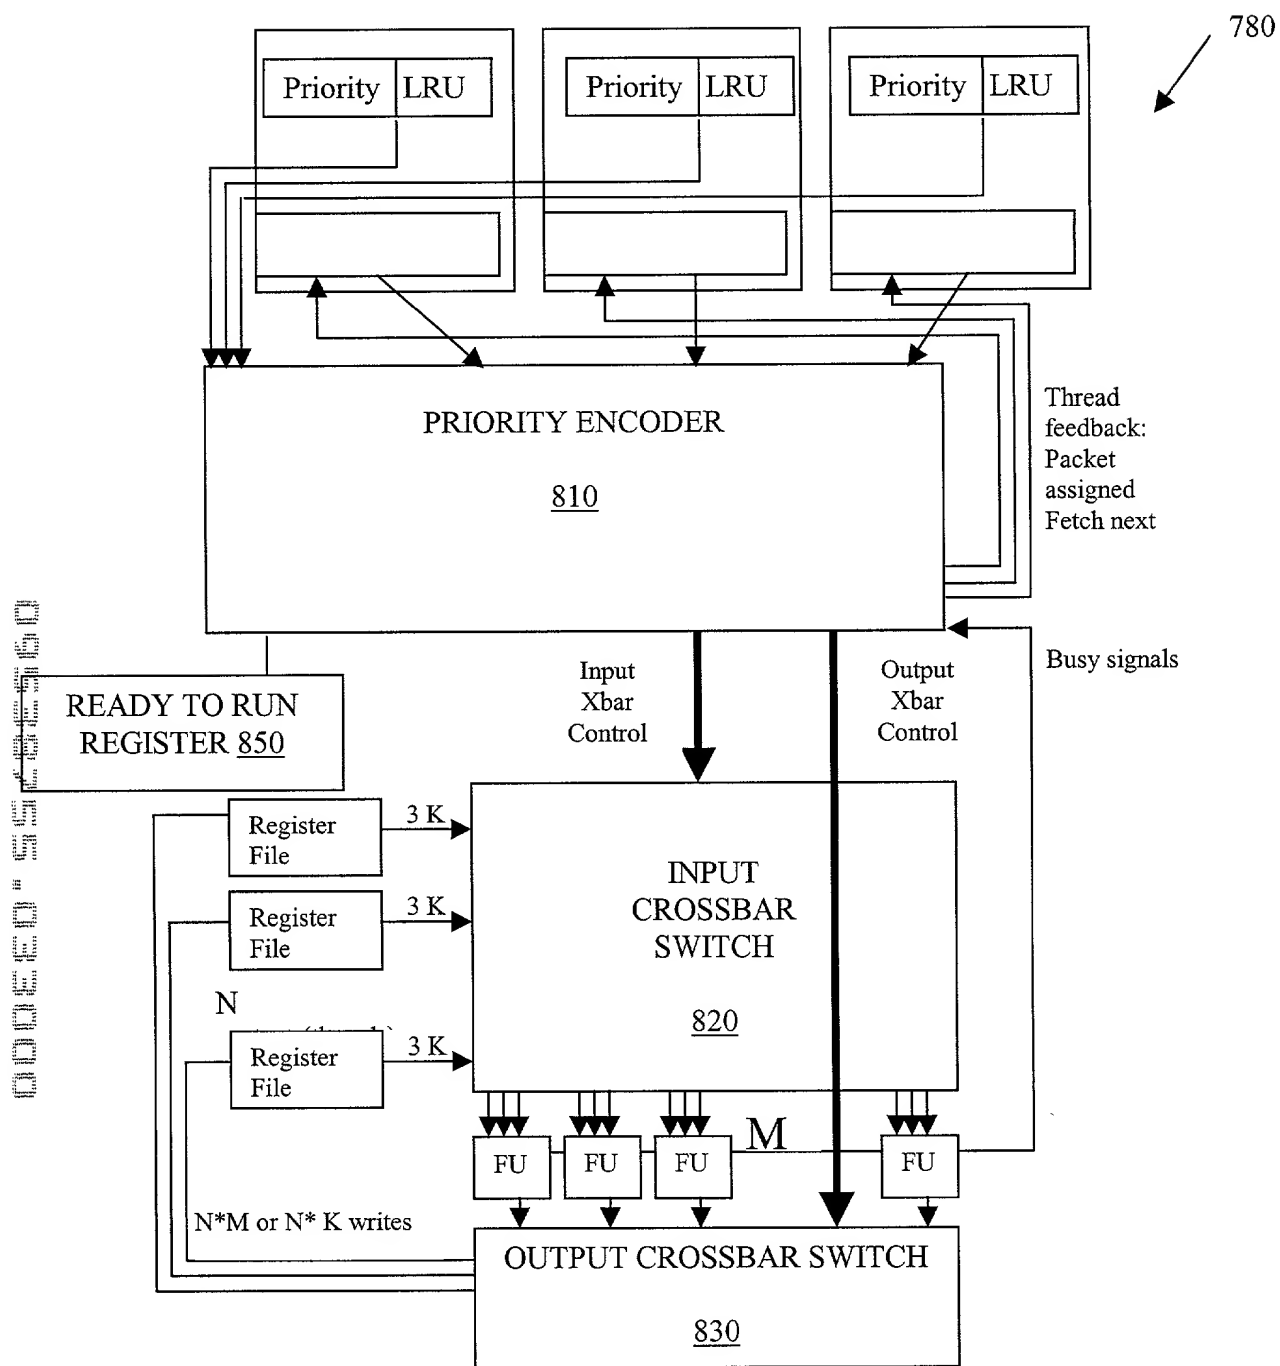


FIG. 8

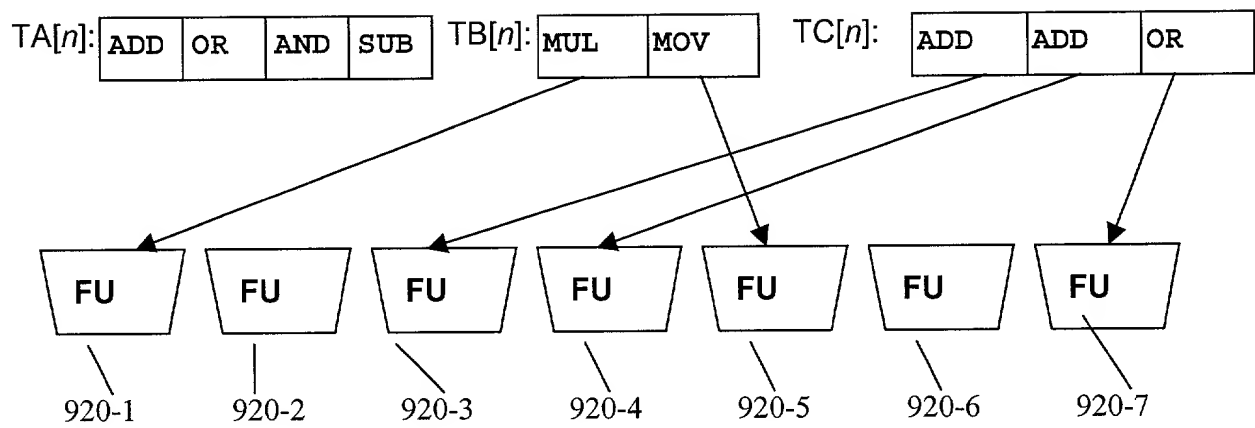


FIG. 9

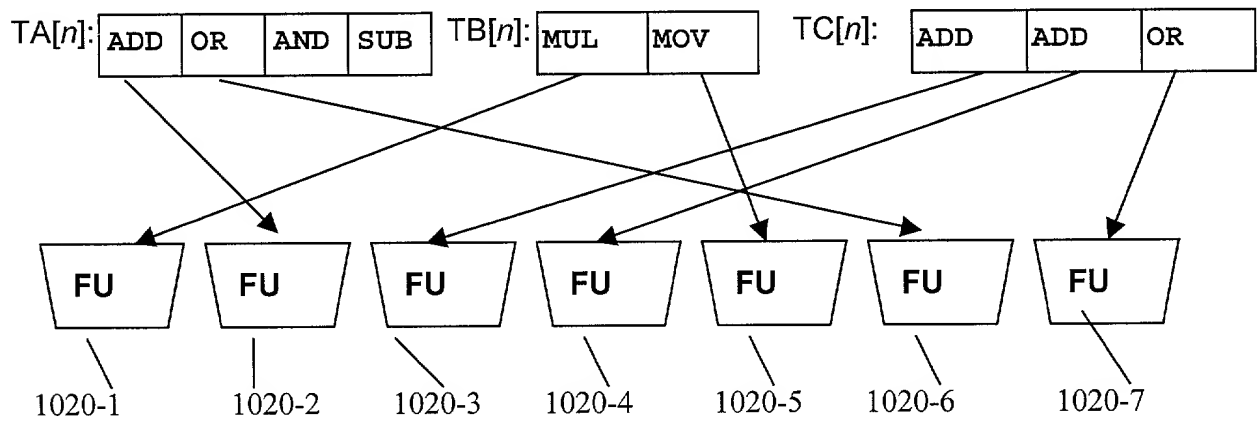


FIG. 10A

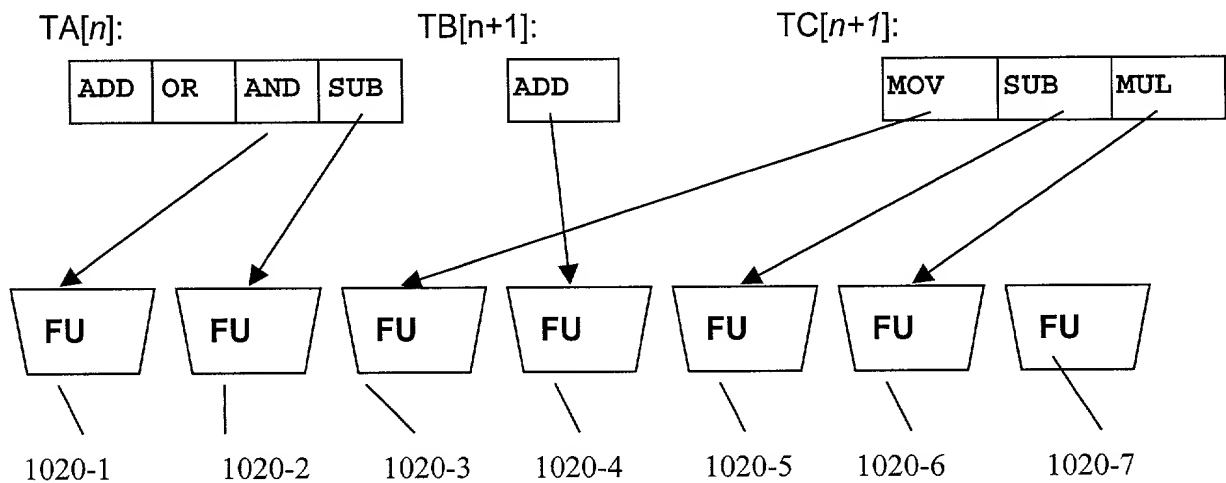


FIG. 10B

Declaration and Power of Attorney

My residence, post office address and citizenship are as stated below next to my name.

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by an amendment, if any, specifically referred to in this oath or declaration.

I hereby claim foreign priority benefits under Title 35, United States Code, 119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

I hereby claim the benefit under Title 35, United States Code, 119(e) of any United States provisional application(s) listed below:

I hereby claim the benefit under Title 35, United States Code, 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, 112, I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

[illegible][illegible]

I hereby appoint the attorneys on ATTACHMENT A as associate attorneys in the aforementioned application, with full power solely to prosecute said application, to make alterations and amendments therein, to receive the patent, and to transact all business in the Patent and Trademark Office connected with the prosecution of said application. No other powers are granted to such associate attorneys and such associate attorneys are specifically denied any power of substitution or revocation.

Full name of 1st joint inventor: Alan David Berenbaum

Inventor's signature Alan David Berenbaum Date 3/29/00

Residence: New York City, Manhattan County, New York

Citizenship: United States of America

Post Office Address: 37 West 12th Street, Apt 9J
New York, NY 10011

Full name of 2nd joint inventor: Nevin Heintze

Inventor's signature Nevin Heintze Date 29/MAR/00

Residence: Morristown, Morris County, New Jersey

Citizenship: Australia

Post Office Address: 301 Mount Kemble Avenue
Morristown, NJ 07960

Full name of 3rd joint inventor: Tor E. Jeremiassen


Inventor's signature Tor E. Jeremiassen Date 3-29-00

Residence: Somerset, Somerset County, New Jersey

Citizenship: Norway

Post Office Address: 81 Oswestry Way
Somerset, NJ 08873

Full name of 4th joint inventor: Stefanos Kaxiras

Inventor's signature  Date 03/29/2000

Residence: Jersey City, Union County, New Jersey

Citizenship: Greece

Post Office Address: 444 Washington Blvd., Apt. 2119
Jersey City, New Jersey

ATTACHMENT A

Attorney Names:	<u>Joseph B. Ryan</u>	Reg. Nos.:	<u>37,922</u>
	<u>Kevin M. Mason</u>		<u>36,597</u>
	<u>William E. Lewis</u>		<u>39,274</u>

Telephone calls should be made to Kevin M. Mason at:

Phone No.: (203) 255-6560

Fax No.: (203) 255-6570

All written communications are to be addressed to:

Ryan & Mason, L.L.P.
90 Forest Avenue
Locust Valley, NY 11560